


## Hybrid Petacomputing Meets Cosmology: The Roadrunner Universe Project

Salman Habib, T-2; Adrian Pope, ISR-1/T-2/CCS-6; Zarija Lukić, T-2; David Daniel, CCS-1; Patricia Fasel, CCS-3; Katrin Heitmann, ISR-1; Nehal Desai, CCS-1 and Aerospace; Chung-Hsing Hsu, CCS-1 and ORNL; Lee Ankeny, HPC-1; Graham Mark, CCS-3; Suman Bhattacharya, T-2; James Ahrens, CCS-1

**Fig. 1. Dark matter halo visualization from a fly-by zoom-in view of a large cosmological simulation with MC<sup>3</sup> (visualization by J. Woodring, CCS-1). Galaxies inherit their clustering properties from those of the host dark matter halos in which they reside.**



Over the last two decades, critical observational advances in large-volume sky surveys carried out over a wide range of wavelengths, as well as over short time cadences, have revolutionized cosmology. Computational cosmology has emerged as an essential resource for providing detailed predictions for these observations, essential data for assisting in their design, and sophisticated tools for interpreting the final results.

Results from cosmological surveys have cemented a cross-validated cosmological Standard Model, presenting a comprehensive picture of the evolutionary history of the Universe and its constituents: 23% in dark matter, which only interacts gravitationally (and a large fraction of which is in localized clumps called halos), and 72% in a smooth dark energy component that is described by a cosmological constant, adiabatic Gaussian random initial density fluctuations, and flat spatial geometry [1]. Although this result is a great triumph, it has exposed some of the biggest puzzles in physical science: What is dark


matter? Why is the expansion of the Universe accelerating? Does general relativity need to be modified? What is the origin of primordial fluctuations?

To investigate these questions, the observational state of the art is rapidly advancing; surveys now coming online and within the next decade represent an

improvement in capability by roughly two orders of magnitude, translating into a determination of certain cosmological parameters at the 1% level. Remarkable as this is, the effort will only come to fruition if the accuracy of the underlying theory can be controlled to the sub-percent level. This severely demanding task will push the boundaries of computing for the foreseeable future.

With one or two exceptions, all cosmological probes are based on statistical studies of fluctuations, whether temperature in the case of the cosmic microwave background (CMB), or of the mass distribution in the case of weak gravitational lensing. The fundamental task for theory is to produce accurate predictions for the CMB, for the mass density and velocity field (along with associated statistical measures such as the fluctuation power spectrum), and for tracers of mass and velocity, such as galaxies and galaxy clusters. Cosmological information on large spatial scales is relatively easier to interpret because on these scales the physics is essentially linear—however, observable quantities can have a large variance due to finite sampling limitations. At smaller, more nonlinear scales, the statistical limitations can be essentially removed, but modeling becomes significantly more complicated, so keeping the associated systematic errors in check is a difficult challenge. There is currently no alternative to precision simulation as the theoretical tool of choice for dealing with the nonlinear regime of structure formation.

Structure formation in the Universe is driven primarily by the gravitational instability. Initial density perturbations collapse and merge in a hierarchical fashion to form dark matter halos within a global *cosmic web* structure (Fig. 1). On scales smaller than several megaparsecs (Mpc) (1 parsec = 3.26 light-years), baryonic matter collects in halos, eventually forming stars and galaxies. The collisionless evolution of matter subject only to gravity is described by the Vlasov-Poisson equation in an expanding Universe, which can be solved in detail only by N-body techniques. Next-generation surveys demand simulations with multi-gigaparsec (Gpc) box-sizes and particle counts in the  $10^{11-12}$  range, all with approximately kiloparsec (kpc) force resolution (a force dynamic range of  $10^6$ ). An overall two to three orders of magnitude improvement in



**Fig. 2.** Scaling of wall-clock time per particle per step as a function of the number of processors. The total problem size is scaled roughly with the number of processors (weak scaling).

throughput over the current state of the art turns out to be the minimal requirement.

To meet the challenge of next-generation simulations, the Roadrunner Universe (RRU) project at LANL, using the IBM Cell Broadband Engine (Cell BE), has developed the Mesh-based Cosmology Code on the Cell BE (MC<sup>3</sup>), the first hybrid petascale cosmology code. The MC<sup>3</sup> algorithm splits the inter-particle force problem into two parts, a medium-resolution solver based on fast Fourier transforms (FFTs), augmented by a direct particle-particle short-range solver. The biggest FFT provides up to four orders of magnitude of dynamic range, the remaining factor of 10–100 coming from the short-range force evaluations carried out on the Cell BE processors of Roadrunner. The MC<sup>3</sup> algorithms match to the machine architecture, minimizing data transfer through the narrow communication pipe between the central processing units (CPU) and the Cell BE. The global philosophy is to explicitly sacrifice memory and in-place computation in order to minimize communication and simplify communication patterns.

Our approach has two key aspects: 1) reduction of particle communication across the Cell layer using particle overloading, a 26-neighboring processor-mirrored particle cache, and 2) application of digital filtering and differencing in the spectral domain, allowing simplified computations at the Cell BE layer [2]. While MC<sup>3</sup> was originally developed for Roadrunner, the computational strategy can easily be modified for other petascale platforms because the short-range and long-range force solvers are effectively decoupled, allowing for independent optimizations. As a result of its design, MC<sup>3</sup> possesses excellent parallel scaling properties (Fig. 2). Since the size of the FFT is quite manageable for current and next-generation supercomputers, MC<sup>3</sup> will scale to the largest problems that can be run on these machines. Additionally, the MC<sup>3</sup> implementation is likely to evolve smoothly as future architectures add further layers of memory and computational hierarchies.

The first scientific application of MC<sup>3</sup> is an analysis of the baryon acoustic oscillation (BAO) signal in the quasar Lyman- $\alpha$  forest [3]. Fixing the BAO scale via cosmological measurements is probably the least systematic-affected technique for investigating the dark energy equation of state. The work involved running nine of some of the largest cosmological simulations ever performed (64 billion particles) with a dynamic range sufficient to resolve the smallest scale of interest (the Jeans scale), as well as enough volume to realistically capture the depth of the Baryon Oscillation Spectroscopic Survey (BOSS), the key cosmological component of the Sloan Digital Sky Survey III. BOSS aims at a percent-level measurement of the BAO scale.

An extensive cosmological simulation program is planned for MC<sup>3</sup>, including applications to next-generation surveys such as the Large Synoptic Survey Telescope (LSST) project along with a variety of physics extensions, such as self-consistent dark energy models, modified gravity, non-Gaussian initial perturbations, and gas physics.

[1] E. Komatsu et al., *Astrophys. J. Supp.* **180**, 330 (2009).

[2] A. Pope et al., *Comp. Sci. Eng.* (to appear).

[3] M. White et al., *Astrophys. J.* arXiv:0911.5341 (to appear).

#### Funding Acknowledgments

LANL Directed Research and Development Program

**For more information contact**  
**Salman Habib** at  
**habib@lanl.gov**.